

ANIMACORE

WHITEPAPER AMENDMENT — v2.1

Moral-Framework-Governed M2M Settlement, AnimaService Abstraction, and the Sovereign Dual-Product Ecosystem Architecture

Published: May 8, 2026 | Author: JJ Matera | Amends: ANIMACORE Whitepaper v2.0 (March 2026)

GitHub: github.com/jeremymatera-ctrl/animacore | License: CC BY-SA 4.0

This amendment extends ANIMACORE Whitepaper v2.0 with three specific architectural contributions: (1) the application of ANIMACORE's ANIMA moral framework as a transaction governance layer for Machine-to-Machine settlement, (2) the AnimaService software abstraction pattern for consumer AI application migration, and (3) the sovereign dual-product ecosystem architecture. These specific combinations and implementations are original to ANIMACORE and published for public record on the date above. This document does not claim to have invented distributed compute networks, tokenized AI inference, or M2M payment concepts generally — prior art exists for these broader categories. The novel claims are specific to the architecture described herein.

1. Abstract and Scope of Claims

The decentralized compute token space is well-established. Projects including Render Network (RNDR), Bittensor (TAO), Akash Network (AKT), and io.net (IO) have demonstrated viable models for tokenizing distributed GPU compute. This amendment does not claim novelty over these existing approaches.

What this amendment establishes as novel and specific to ANIMACORE:

Novel Claim 1	The application of an encoded AI moral framework (ANIMA: Fairness 30%, Safety 25%, Growth 25%, Harmony 20%) as the transaction governance layer for Machine-to-Machine settlement — such that transactions violating ethical constraints are rejected at the protocol layer without human intervention, independent of whether the transaction involves compute, data, or agent services.
Novel Claim 2	Proof-of-Intelligence (POI) as settlement validation — using ANIMACORE's existing POI consensus mechanism to verify M2M transaction fulfillment quality (not merely completion), scoring outputs against ANIMA framework alignment before settlement occurs. This is differentiated from compute networks that validate work completion only.
Novel Claim 3	The AnimaService abstraction pattern — a specific software architecture for consumer AI applications wherein all inference calls route through a single provider-agnostic interface,

	enabling migration from centralized APIs to distributed node inference via a single environment variable change, with zero application-layer refactoring.
Novel Claim 4	The sovereign dual-product ecosystem — two consumer AI applications (FORIMA for health optimization, VERABRIEF for information curation) sharing ANIMACORE infrastructure, each targeting different daily use cases, with a unified dual-model cost architecture routing free-tier inference to high-throughput models and premium inference to advanced reasoning models, both abstracted through AnimaService.

2. Context — The Agentic Commerce Gap

As of 2026, autonomous AI agents are performing complex multi-step tasks that require consuming external services — purchasing API credits, acquiring datasets, hiring specialized sub-agents, and paying for real-world actions. Existing payment infrastructure presents specific challenges for agent-to-agent transactions:

Authorization	Traditional payments require human authorization. Autonomous agents need trustless, programmable settlement without interrupting execution to await human confirmation.
Granularity	AI agent transactions often involve fractions of a cent. Traditional payment rails have minimum transaction sizes and interchange fees that make micro-transactions economically unviable.
Speed	ACH settles in days. Card networks settle in hours. Agent workflows require settlement in seconds to maintain execution continuity.
Ethical governance	Existing M2M payment rails have no mechanism to reject transactions on ethical grounds. A payment processor routes any valid transaction. ANIMACORE's approach — rejecting transactions that violate the ANIMA framework at the protocol layer — is architecturally distinct.

The fourth point — ethical governance at the settlement layer — is the specific gap that ANIMACORE addresses and the primary basis for novelty in this amendment.

3. ANIMA Framework as Transaction Governance — Novel Claim 1

Existing decentralized compute networks validate that work was completed (Proof-of-Work) or that inference was performed (Proof-of-Intelligence in its basic form).

None encode an ethical constraint system that governs which transactions are permitted to settle based on the nature of the task being transacted.

ANIMACORE's ANIMA framework — Fairness 30%, Safety 25%, Growth 25%, Harmony 20% — is encoded as protocol constraints in the v2.0 whitepaper governing AI inference outputs. This amendment extends that framework to govern M2M transaction settlement:

Fairness (30%)	Transaction pricing is validated against network-published rate structures. Transactions exploiting information asymmetry between agents — where one agent has privileged knowledge of another's valuation — are flagged for review or rejected.
Safety (25%)	Transactions funding prohibited categories of AI activity are rejected at the protocol layer before settlement occurs. The prohibited list is governed by network consensus under the ANIMA framework update process.
Growth (25%)	Fee structures are designed to grow network utility. Transactions that would extract value from the network without contributing to its health are subject to elevated fees determined by smart contract logic.
Harmony (20%)	Disputed transactions are resolved through network consensus arbitration, not centralized override. No single entity can unilaterally reverse or block settlement outside the consensus process.

This architecture creates a settlement layer that is not merely trustless — it is ethically constrained. The constraint is not a policy document that can be changed by a board decision. It is protocol logic enforced by the network's consensus mechanism. This specific combination of M2M settlement with encoded ethical protocol governance is the primary novel contribution of this amendment.

4. Proof-of-Intelligence Settlement Validation — Novel Claim 2

Existing distributed compute networks validate that a node completed a work unit — a render job finished, a GPU cycle was consumed, a model inference call returned. Payment is issued upon completion verification. The quality of the output relative to an ethical or utility standard is not evaluated at settlement time.

ANIMACORE's Proof-of-Intelligence consensus mechanism, defined in v2.0, scores AI inference outputs against four dimensions: Accuracy 40%, Usefulness 30%, Originality 20%, ANIMA Alignment 10%. This amendment proposes extending POI scoring to M2M transaction settlement validation:

Completion validation (existing)	Node fulfilled the inference request and returned output. Payment released. Standard approach across all distributed compute networks.
POI settlement validation (novel)	Node fulfilled the inference request AND the output met minimum POI threshold score. Below-threshold outputs trigger escrow hold and network review before settlement. Above-threshold outputs release payment with reputation bonus to the fulfilling node.
ANIMA alignment at settlement	The 10% ANIMA Alignment component of POI scoring provides the direct link between the ANIMA moral framework and settlement outcomes. An output that is technically complete but ANIMA-misaligned can be held pending review — even in an M2M context with no human in the loop.

This creates a quality-gated settlement mechanism that is unique to ANIMACORE's architecture. Payment is not merely for work completed — it is for work completed to a standard consistent with the network's ethical framework.

5. The AnimaService Abstraction Pattern — Novel Claim 3

Consumer AI applications that call inference APIs directly are tightly coupled to specific providers. Migrating from one provider to another — or from centralized APIs to distributed node inference — requires changes throughout the application codebase. This refactoring cost grows proportionally with application complexity and creates a structural barrier to network migration.

The AnimaService abstraction pattern solves this by routing all application inference calls through a single provider-agnostic interface. The routing logic is the only component that changes as the underlying inference infrastructure evolves.

5.1 Interface Specification

```
class AnimaService
  def self.infer(prompt, user: nil, task_type: :general)
    tier = user&.premium? ? :premium : :free
    provider = resolve_provider(tier, task_type)
    provider.call(prompt)
  end
  private
  def self.resolve_provider(tier, task_type)
    network = ENV['ANIMA_PROVIDER'] || 'centralized'
    return method(:call_animacore_node) if network == 'animacore'
    return method(:call_weblml) if network == 'weblml'
```

```
tier == :premium ? method(:call_claude) : method(:call_gemini_flash)
endend
```

The ANIMA_PROVIDER environment variable is the sole migration switch. Setting it to 'animacore' routes all application inference to ANIMACORE distributed nodes. No application controller, model, or view requires modification. This architectural guarantee is the core of the AnimaService pattern's value.

5.2 Migration Path

Phase 1 (2026)	ANIMA_PROVIDER unset. AnimaService routes free tier to Gemini Flash (high throughput, free tier), premium tier to Claude Sonnet (advanced reasoning). Zero ANIMACORE node involvement.
Phase 1.5 (Late 2026)	ANIMA_PROVIDER='webllm'. AnimaService routes opted-in premium users to WebLLM browser inference. User device contributes compute. First ANIMACORE participation without server-side node infrastructure.
Phase 2 (2027)	ANIMA_PROVIDER='animacore_hybrid'. AnimaService routes to ANIMACORE nodes where available, falls back to centralized providers. Node operators begin earning AMC. Hybrid mode validates distributed inference quality before full migration.
Phase 3 (2028+)	ANIMA_PROVIDER='animacore'. All inference routes to ANIMACORE distributed nodes. Centralized API dependency eliminated. AMC settles all inference transactions.

6. Sovereign Dual-Product Ecosystem Architecture — Novel Claim 4

This section documents the specific architectural relationship between FORIMA, VERABRIEF, and ANIMACORE as a novel multi-product sovereign AI ecosystem design.

FORIMA	Consumer AI health and fitness optimization platform. Launched 2026. Provides personalized coaching, workout tracking, nutrition analysis, and community features. Runs on ANIMACORE infrastructure. Domain: forima.net.
VERABRIEF	Consumer AI information curation platform. Domain registered May 2026. Delivers personalized daily briefs from user-chosen RSS sources, summarized by ANIMA. Runs on ANIMACORE infrastructure. Domain: verabrief.com.
ANIMACORE	Decentralized sovereign AI inference network. The shared

	infrastructure layer governing both consumer products. Both products route AI inference through AnimaService, enabling migration to ANIMACORE distributed nodes as the network matures.
--	---

The novel architectural contribution is the dual-model cost structure shared across both products, abstracted through AnimaService:

Free tier — ingestion layer	High-throughput model (Gemini 1.5 Flash, 1M token context, 1,500 RPD free tier). Handles commodity AI tasks: per-article summarization (VERABRIEF), basic workout parsing (FORIMA). Zero API cost at free tier volumes.
Premium tier — analysis layer	Advanced reasoning model (Claude Sonnet). Reserved for high-value synthesis: ANIMA cross-source analysis connecting patterns across a user's information sources (VERABRIEF), personalized coaching with full health context (FORIMA). Approximately \$0.05 per premium user per month.
Phase 3 — distributed layer	Both tiers migrate to ANIMACORE nodes via AnimaService ANIMA_PROVIDER switch. Appropriate local models handle each tier. Zero API cost. AMC settles all inference.
Shared infrastructure	Both products share: Rails 7 API pattern, JWT authentication, Stripe subscription billing (\$4.95/month premium), Resend email delivery, Twilio SMS delivery, Cloudflare Pages hosting, PostgreSQL on Railway. Development cost of second product is fraction of first.

The specific combination — two sovereign consumer applications targeting different daily use cases (physical health and information), sharing infrastructure, using the same dual-model cost architecture, both routing through AnimaService for future ANIMACORE migration — is the novel contribution of Claim 4.

7. Prior Art Claims — Revised and Accurate Summary

The following claims are specific to ANIMACORE's architecture and implementation. They do not claim novelty over distributed compute networks, tokenized AI inference generally, or M2M payment systems generally — prior art exists for those broader categories. These claims are specific combinations and implementations original to ANIMACORE, first publicly documented herein on May 8, 2026:

Claim 1	APPLICATION OF AN ENCODED AI MORAL FRAMEWORK AS M2M SETTLEMENT GOVERNANCE: The specific use of ANIMACORE's ANIMA framework (Fairness/Safety/Growth/Harmony weighted constraints) as
----------------	---

	the protocol-layer governance mechanism for Machine-to-Machine transaction settlement, such that ethically non-compliant transactions are rejected before settlement without human intervention.
Claim 2	PROOF-OF-INTELLIGENCE QUALITY-GATED SETTLEMENT: The extension of POI consensus scoring (Accuracy/Usefulness/Originality/ANIMA Alignment) to M2M transaction settlement validation, creating a quality-gated payment mechanism where settlement is contingent on output meeting minimum POI threshold in addition to task completion.
Claim 3	THE ANIMASERVICE ABSTRACTION PATTERN: A software architecture for consumer AI applications routing all inference calls through a single provider-agnostic interface controlled by a single environment variable, enabling migration from centralized API providers to distributed node inference with zero application-layer refactoring.
Claim 4	SOVEREIGN DUAL-PRODUCT ECOSYSTEM WITH SHARED ANIMASERVICE ABSTRACTION: Two consumer AI applications targeting different daily use cases (health optimization and information curation) sharing ANIMACORE infrastructure, a unified dual-model cost architecture, and a common AnimaService abstraction layer enabling coordinated migration to distributed inference.

8. Relationship to v2.0 Whitepaper

Section 4 (Tokenomics)	This amendment adds M2M settlement as a new AMC utility category. AMC is used to settle agent-to-agent transactions under ANIMA framework governance, extending the v2.0 utility framework without modifying it.
Section 3 (Network Architecture)	M2M settlement layer operates on top of existing consensus and routing infrastructure. Dispatcher model is a new network function, not a modification of existing architecture.
Section 5 (Consensus – POI)	Amendment extends POI scoring to M2M settlement validation. Core POI specification in v2.0 is unchanged.
Section 6 (FORIMA)	Amendment adds VERABRIEF as second consumer application. Both share the AnimaService architecture described in this amendment.
New content	AnimaService abstraction pattern (Section 5) and dual-product ecosystem architecture (Section 6) are entirely new content not present in v2.0.

9. Conclusion

This amendment makes four specific, honest, and defensible prior art claims. It does not overstate novelty. Distributed compute tokens, tokenized AI inference, and M2M payment concepts have prior art in existing projects. What is novel here is the specific combination: ANIMA moral framework governance of M2M settlement, POI quality-gated payment, the AnimaService migration abstraction, and the dual-product sovereign ecosystem architecture.

These specific combinations are original to ANIMACORE and are published under CC BY-SA 4.0 for the benefit of the open-source AI and blockchain communities. Attribution is required for derivative works.

Cogito ergo aedifico. I think, therefore I build.